arXiv:1409.4132v1 [cs.GT] 15 Sep 2014

Equilibria of Plurality Voting: Lazy and Truth-biased Voters

Edith Elkind* Evangelos Markakis† Svetlana Obraztsova‡
Piotr Skowron§

Abstract

We present a systematic study of Plurality elections with strategic voters who, in addition to having preferences over election winners, have secondary preferences, which govern their behavior when their vote cannot affect the election outcome. Specifically, we study two models that have been recently considered in the literature [6, 17]: *lazy* voters, who prefer to abstain when they are not pivotal, and *truth-biased* voters, who prefer to vote truthfully when they are not pivotal. We extend prior work by investigating the behavior of both lazy and truth-biased voters under different tie-breaking rules (lexicographic rule, random voter rule, random candidate rule). Two of these six combinations of secondary preferences and a tie-breaking rule have been studied in prior work. In order to understand the impact of different secondary preferences and tie-breaking rules on the election outcomes, we study the remaining four combinations. We characterize pure Nash equilibria (PNE) of the resulting strategic games and study the complexity of related computational problems. Our results extend to settings where some of the voters may be non-strategic.

1 Introduction

Plurality voting is a popular tool for collective decision-making in many domains, including both human societies and multiagent systems. Under this voting rule, each voter is supposed to vote for her most favorite candidate (or abstain); the winner is then the candidate that receives the highest number of votes. If several candidates have the highest score, the winner is chosen among them using a *tie-breaking rule*; popular tie-breaking rules include the *lexicographic rule*, which imposes a fixed priority order over the candidates; the *random candidate rule*, which picks one of the tied candidates uniformly at random; and the *random voter rule*, which picks the winner among the tied candidates according to the preferences of a randomly chosen voter.

*University of Oxford
†Athens University of Economics and Business
‡National Technical University of Athens
§University of Warsaw

In practice, voters are often *strategic*, i.e., they may vote non-truthfully if they can benefit from doing so. In that case, an election can be viewed as a game, where the voters are the players, and each player's space of actions includes voting for any candidate or abstaining. For deterministic rules (such as Plurality with lexicographic tie-breaking), the behavior of strategic voters is determined by their preference ordering, i.e., a ranking of the candidates, whereas for randomized rules a common approach is to specify utility functions for the voters; i.e., the voters are assumed to maximize their *expected utility* under the lottery induced by tie-breaking. The outcome of the election can then be identified with a pure Nash equilibrium (PNE) of the resulting game.

However, for the Plurality voting game with 3 or more voters, this approach fails to provide a useful prediction of voting behavior: for each candidate c there is a PNE where c is the unique winner, irrespective of the voters' preferences. Indeed, if there are at least 3 voters, the situation where all of them vote for c is a PNE, as no voter can unilaterally change the election outcome. However, such equilibria may disappear if we use a more refined model of voters' preferences that captures additional aspects of their decision-making. For instance, in practice, if a voter feels that her vote is unlikely to have any effect on the election outcome, she may decide to abstain from the election. Also, voters may be averse to lying about their preferences, in which case they can be expected to vote for their top candidate unless there is a clear strategic reason to vote for someone else. By taking into account these aspects of voters' preferences, we obtain a more faithful model of their behavior.

The problem of characterizing and computing the equilibria of Plurality voting, both for "lazy" voters (i.e., ones who prefer to abstain when they are not pivotal) and for "truth-biased" voters (ones who prefer to vote truthfully when they are not pivotal), has recently received a considerable amount of attention. However, it is difficult to compare the existing results, since they rely on different tie-breaking rules. In particular, [6], who study lazy voters, use the random candidate tie-breaking rule, and [17] consider truth-biased voters and the lexicographic tie-breaking rule. Thus, it is not clear whether the differences between the results in these papers can be attributed to voters' secondary preferences or to the tie-breaking rule.

The primary goal of our paper is to tease out the effects of different features of these models, by systematically considering various combinations of secondary preferences and tie-breaking rules. We consider two types of secondary preferences (lazy voters and truth-biased voters) and three tie-breaking rules (the lexicographic rule, the random voter rule, and the random candidate rule); while two of these combinations have been studied earlier by Desmedt and Elkind [6] and Obraztsova et al. [17], to the best of our knowledge, the remaining four possibilities have not been considered before. For each of the new scenarios, we characterize the set of PNE for the resulting game; in doing so, we also fill in a gap in the characterization of Desmedt and Elkind for lazy voters and random candidate tie-breaking. We then consider the problems of deciding whether a given game admits a PNE and whether a given candidate can be a co-winner/unique winner in some PNE of a given game. For all settings we consider, we determine the computational complexity of each of these problems, classifying them as either polynomial-time solvable or NP-complete. We use our characterization results to analyze the impact of various features of our models on the election outcomes. Finally, we extend our results to the setting where some of the voters may be *principled*,

i.e., are guaranteed to vote truthfully.

Related Work

Equilibria of Plurality voting have been investigated by a number of researchers, starting with [9]. However, most of the earlier works either consider solution concepts other than pure Nash equilibria, such as iterative elimination of dominated strategies [13, 7], or assume that voters have incomplete information about each others' preferences [14]. Both types of secondary preferences (lazy voters and truth-biased voters) appear in the social choice literature, see, respectively, [2, 3, 19] and [8, 11]. In computational social choice, truth-biased voters have been considered by Meir et al. [12] in the context of dynamics of Plurality voting; subsequently, Plurality elections with truth-biased voters have been investigated empirically by Thompson et al. [20] and theoretically by Obraztsova et al. [17]. To the best of our knowledge, the only paper to study computational aspects of Plurality voting with lazy voters is that of Desmedt and Elkind [6].

Our approach to tie-breaking is well-grounded in existing works. Lexicographic tie-breaking is standard in the computational social choice literature. The random candidate rule has been discussed by Desmedt and Elkind [6], and, more recently, by Obraztsova, Elkind and Hazon [16] and Obraztsova and Elkind [15]. The random voter rule is used to break ties under the Schulze method [18]; complexity of manipulation under this tie-breaking rule has been studied by Aziz et al. [1].

2 Preliminaries

For any positive integer t, we denote the set $\{1, \ldots, t\}$ by $[t]$. We consider elections with a set of *voters* $N = [n]$ and a set of *alternatives*, or *candidates*, $C = \{c_1, \ldots c_m\}$. Each voter is associated with a *preference order*, i.e., a strict linear order over C; we denote the preference order of voter i by $\succ_i$. The list $(\succ_1, \ldots, \succ_n)$ is called a *preference profile*. For each $i \in N$, we set a_i to be the top choice of voter i, and let $\mathbf{a} = (a_1, \ldots, a_n)$. Given two disjoint sets of candidates X, Y and a preference order $\succ$, we write $X \succ Y$ if in $\succ$ all candidates from X are ranked above all candidates from Y.

We also assume that each voter $i \in N$ is endowed with a *utility function* $u_i : C \to \mathbb{N}$; $u_i(c_j)$ is the utility derived by voter i if c_j is the unique election winner. We require that $u_i(c) \neq u_i(c')$ for all $i \in N$ and all $c, c' \in C$ such that $c \neq c'$. The vector $\mathbf{u} = (u_1, \ldots, u_n)$ is called the *utility profile*. Voters' preference orders and utility functions are assumed to be consistent, i.e., for each $i \in N$ and every pair of candidates $c, c' \in C$ we have $c \succ_i c'$ if and only if $u_i(c) > u_i(c')$; when this is the case, we will also say that $\succ_i$ is *induced* by u_i. Sometimes, instead of specifying preference orders explicitly, we will specify the utility functions only, and assume that voters' preference orders are induced by their utility functions; on other occasions, it will be convenient to reason in terms of preference orders.

A *lottery* over C is a vector $\mathbf{p} = (p_1, \ldots, p_m)$ with $p_j \geq 0$ for all $j \in [m]$ and $\sum_{j \in [m]} p_j = 1$. The value p_j is the probability assigned to candidate c_j. The *expected utility* of a voter $i \in N$ from a lottery $\mathbf{p}$ is given by $\sum_{j \in [m]} u_i(c_j) p_j$.

In this paper we consider Plurality elections. In such elections each voter $i \in N$ submits a *vote*, or *ballot*, $b_i \in C \cup \{\perp\}$; if $b_i = \perp$, voter i is said to *abstain*. The list of all votes $\mathbf{b} = (b_1, \ldots, b_n)$ is also called a *ballot vector*. We say that a ballot vector is *trivial* if $b_i = \perp$ for all $i \in N$. Given a ballot vector $\mathbf{b}$ and a ballot b', we write $(\mathbf{b}_{-i}, b')$ to denote the ballot vector obtained from $\mathbf{b}$ by replacing b_i with b'. The *score* of an alternative c_j in an election with ballot vector $\mathbf{b}$ is given by $\mathrm{sc}(c_j, \mathbf{b}) = |\{i \in N \mid b_i = c_j\}|$. Given a ballot vector $\mathbf{b}$, we set $M(\mathbf{b}) = \max_{c \in C} \mathrm{sc}(c, \mathbf{b})$ and let $W(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b})\}$, $H(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 1\}$, $H'(\mathbf{b}) = \{c \in C \mid \mathrm{sc}(c, \mathbf{b}) = M(\mathbf{b}) - 2\}$. The set $W(\mathbf{b})$ is called the *winning set*. Note that if $\mathbf{b}$ is trivial then $W(\mathbf{b}) = C$. If $|W(\mathbf{b})| = 1$ then the unique candidate in $W(\mathbf{b})$ is declared to be the winner. Otherwise, the winner is selected from $W(\mathbf{b})$ according to one of the following tie-breaking rules.

(1) Under the *lexicographic rule* R^L, the winner is the candidate $c_j \in W(\mathbf{b})$ such that $j \leq k$ for all $c_k \in W(\mathbf{b})$.

(2) Under the *random candidate rule* R^C, the winner is chosen from $W(\mathbf{b})$ uniformly at random.

(3) Under the *random voter rule* R^V, we select a voter from N uniformly at random; if she has voted for a candidate in $W(\mathbf{b})$, we output this candidate, otherwise we ask this voter to report her most preferred candidate in $W(\mathbf{b})$, and output the answer. This additional elicitation step may appear difficult to implement in practice; fortunately, we can show that, in equilibrium it is almost never necessary.

Thus, the outcome of an election is a lottery over C; however, for R^L this lottery is degenerate, i.e., it always assigns the entire probability mass to a single candidate. For each $X \in \{L, C, V\}$ and each ballot vector $\mathbf{b}$, let $\mathbf{p}^X(\mathbf{b})$ denote the lottery that corresponds to applying R^X to the set $W(\mathbf{b})$. Note also that for every $c_j \in C$ it holds that if $p_j^C(\mathbf{b}) \neq 0$ then $p_j^C(\mathbf{b}) \geq \frac{1}{m}$; similarly, if $p_j^V(\mathbf{b}) \neq 0$ then $p_j^V(\mathbf{b}) \geq \frac{1}{n}$.

In what follows, we consider *lazy* voters, who prefer to abstain when their vote has no effect on the election outcome, and *truth-biased* voters, who never abstain, but prefer to vote truthfully when their vote has no effect on the election outcome. Formally, pick $\varepsilon < \min\{\frac{1}{m}, \frac{1}{n}\}$, and consider a utility profile $\mathbf{u}$ and a tie-breaking rule $R^X \in \{R^C, R^V, R^L\}$. Then

- if voter i is *lazy*, her utility in an election with ballot vector $\mathbf{b}$ under tie-breaking rule R^X is given by

$$U_i(\mathbf{b}) = \begin{cases} \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j) & \text{if } b_i \in C, \\ \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j) + \varepsilon & \text{if } b_i = \perp. \end{cases}$$

- if voter i is *truth-biased*, her utility in an election with ballot vector $\mathbf{b}$ under tie-breaking rule R^X is given by

$$U_i(\mathbf{b}) = \begin{cases} \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j) & \text{if } b_i \in C \setminus \{a_i\}, \\ \sum_{j \in [m]} p_j^X(\mathbf{b}) u_i(c_j) + \varepsilon & \text{if } b_i = a_i, \\ -\infty & \text{if } b_i = \perp. \end{cases}$$

We consider settings where all voters are of the same type, i.e., either all voters are lazy or all voters are truth-biased; we refer to these settings as *lazy* or *truth-biased*, respectively, and denote the former by $\mathcal{L}$ and the latter by $\mathcal{T}$.

In what follows, we consider all possible combinations of settings ($\mathcal{L}$, $\mathcal{T}$) and tie-breaking rules (R^L, R^C, R^V). A combination of a setting $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$, a tie-breaking rule $R \in \{R^L, R^C, R^V\}$ and a utility profile $\mathbf{u}$ induces a strategic game, which we will denote by $(\mathcal{S}, R, \mathbf{u})$: in this game, the players are the voters, the action space of each player is $C \cup \{\perp\}$, and the players' utilities $U_1, \ldots, U_n$ for a vector of actions $\mathbf{b}$ are computed based on the setting and the tie-breaking rule as described above. We say that a ballot vector $\mathbf{b}$ is a *pure Nash equilibrium (PNE)* of the game $(\mathcal{S}, R, \mathbf{u})$ if $U_i(\mathbf{b}) \geq U_i(\mathbf{b}_{-i}, b')$ for every voter $i \in N$ and every $b' \in C \cup \{\perp\}$.

For each setting $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$ and each tie-breaking rule $R \in \{R^L, R^C, R^V\}$, we define three algorithmic problems, which we call $(\mathcal{S}, R)$-EXISTNE, $(\mathcal{S}, R)$-TIENE, and $(\mathcal{S}, R)$-SINGLENE. In each of these problems, we are given a candidate set C, $|C| = m$, a voter set N, $|N| = n$, and a utility vector $\mathbf{u} = (u_1, \ldots, u_n)$, where each u_i is represented by m numbers $u_i(c_1), \ldots, u_i(c_m)$; these numbers are positive integers given in binary. In $(\mathcal{S}, R)$-TIENE and $(\mathcal{S}, R)$-SINGLENE we are also given the name of a target candidate $c_p \in C$. In $(\mathcal{S}, R)$-EXISTNE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE. In $(\mathcal{S}, R)$-TIENE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $|W(\mathbf{b})| > 1$ and $c_p \in W(\mathbf{b})$. In $(\mathcal{S}, R)$-SINGLENE we ask if $(\mathcal{S}, R, \mathbf{u})$ has a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{c_p\}$. Each of these problems is obviously in NP, as we can simply guess an appropriate ballot vector $\mathbf{b}$ and check that it is a PNE.

We omit some of the proofs due to space constraints; these proofs can be found in the supplementary material.

3 Lazy Voters

In this section, we study PNE in Plurality games with lazy voters. The case where the tie-breaking rule is R^C has been analyzed in detail by Desmedt and Elkind [6], albeit for a slightly different model; we complement their results by considering R^L and R^V.

We start by extending a result of Desmedt and Elkind to all three tie-breaking rules considered in this paper.

Proposition 1. *For every $R \in \{R^L, R^C, R^V\}$ and every utility profile $\mathbf{u}$, if a ballot vector $\mathbf{b}$ is a PNE of $(\mathcal{L}, R, \mathbf{u})$ then for every voter $i \in N$ either $b_i = \perp$ or $b_i \in W(\mathbf{b})$. Further, if $|W(\mathbf{b})| = 1$, then there exists exactly one voter $i \in N$ with $b_i \neq \perp$.*

Proof. Suppose that $b_i \notin W(\mathbf{b})$ for some voter $i \in N$. Then if i changes her vote to $\perp$, the set $W(\mathbf{b})$ will not change, so i's utility would improve by ε, a contradiction with $\mathbf{b}$ being a PNE of $(\mathcal{L}, R, \mathbf{u})$. Similarly, suppose that $|W(\mathbf{b})| = 1$ and there are two voters $i, i' \in N$ with $b_i \neq \perp$, $b_{i'} \neq \perp$. It has to be the case that $b_i = b_{i'} = c_j$ for some $c_j \in C$, since otherwise $|W(\mathbf{b})| \geq 1$. But then if voter i changes her vote to $\perp$, c_j will remain the election winner, so i's utility would improve by ε, a contradiction. □

Lexicographic Tie-breaking The scenario where voters are lazy and ties are broken lexicographically turns out to be fairly easy to analyze.

Theorem 1. *For any utility profile* $\mathbf{u}$ *the game* $G = (\mathcal{L}, R^L, \mathbf{u})$ *has the following properties:*

1. *If* $\mathbf{b}$ *is a PNE of* G *then* $|W(\mathbf{b})| \in \{1, m\}$*. Moreover,* $|W(\mathbf{b})| = m$ *if and only if* $\mathbf{b}$ *is the trivial ballot and all voters rank* c_1 *first.*
2. *If* $\mathbf{b}$ *is a PNE of* G *then there exists at most one voter* i *with* $b_i \neq \perp$*.*
3. G *admits a PNE if and only if all voters rank* c_1 *first (in which case* c_1 *is the unique PNE winner) or there exists a candidate* c_j *with* $j > 1$ *such that (i)* $\mathrm{sc}(c_j, \mathbf{a}) > 0$ *and (ii) for every* $k < j$ *it holds that all voters prefer* c_j *to* c_k*. If such a candidate exists, he is unique, and wins in all PNE of* G*.*

The following corollary is directly implied by Theorem 1.

Corollary 1. $(\mathcal{L}, R^L)$-EXISTNE, $(\mathcal{L}, R^L)$-SINGLENE *and* $(\mathcal{L}, R^L)$-TIENE *are in* P.

Remark 1. The reader may observe that, counterintuitively, while the lexicographic tie-breaking rule appears to favor c_1, it is impossible for c_1 to win the election unless he is ranked first by all voters. In contrast, c_2 wins the election as long as he is ranked first by at least one voter and no voter prefers c_1 to c_2. In general, the lexicographic tie-breaking rule favors lower-numbered candidates with the exception of c_1. As for c_1, his presence mostly has a destabilizing effect: if some, but not all voters rank c_1 first, no PNE exists. This phenomenon is an artifact of our treatment of the trivial ballot vector: it disappears if we assume (as Desmedt and Elkind do) that when $\mathbf{b} = (\perp, \ldots, \perp)$ the election is declared invalid and the utility of each voter is $-\infty$: under this assumption c_1 is the unique possible equilibrium winner whenever he is ranked first by at least one voter.

Randomized Tie-breaking We will now consider R^C and R^V. [6] characterize utility profiles that admit a PNE for lazy voters and R^C. However, there is a small difference between our model and that of Desmedt and Elkind: while we assume that the trivial ballot vector results in a tie among all candidates, Desmedt and Elkind assume that in this case the election is canceled and each voter's utility is $-\infty$. Further, the results of Desmedt and Elkind implicitly assume that the number of voters n exceeds the number of candidates m; if this is not the case, Theorem 2 in their paper is incorrect (see Remark 2).

Thus, we will now provide a full characterization of utility profiles $\mathbf{u}$ such that $(\mathcal{L}, R^C, \mathbf{u})$ admits a PNE, and describe the corresponding equilibrium ballot profiles. Our characterization result remains essentially unchanged if we replace R^C with R^V: for almost all utility profiles $\mathbf{u}$ and ballot vectors $\mathbf{b}$ it holds that $\mathbf{b}$ is a PNE of $(\mathcal{L}, R^C, \mathbf{u})$ if and only if it is a PNE of $(\mathcal{L}, R^V, \mathbf{u})$; the only exception is the case of full consensus (all voters rank the same candidate first).

Theorem 2. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C*,* $|C| = m$*, and let* $R \in \{R^C, R^V\}$*. The game* $G = (\mathcal{L}, R, \mathbf{u})$ *admits a PNE if and only if one of the following conditions holds:*

(1) all voters rank some candidate c_j *first;*

(2) each candidate is ranked first by at most one voter, and, moreover, $\frac{1}{n}\sum_{i\in N} u_\ell(a_i) \geq \max_{i\in N\setminus\{\ell\}} u_\ell(a_i)$ *for each* $\ell \in N$.

(3) there exists a set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ *with* $2 \leq k \leq \min(n/2, m)$ *and a partition of the voters into* k *groups* $N_1, \ldots, N_k$ *of size* n/k *each such that for each* $j \in [k]$ *and each* $i \in N_j$ *we have* $c_{\ell_j} \succ_i c$ *for all* $c \in X \setminus \{c_{\ell_j}\}$*, and, moreover,* $\frac{1}{k}\sum_{c\in X} u_i(c) \geq \max_{c\in X\setminus\{c_{\ell_j}\}} u_i(c)$.

Further, if condition (1) holds for some $c_j \in C$*, then if* $R = R^C$ *then for each* $i \in N$ *the game* G *has a PNE where* i *votes for* c_j *and all other voters abstain, whereas if* $R = R^V$ *the game* G *has a PNE where all voters abstain; if condition (2) holds, then* G *has a PNE where each voter votes for her top candidate; and if condition (3) holds for some set* X*, then* G *has a PNE where each voter votes for her favorite candidate in* X*. The game* G *has no other PNE.*

Remark 2. Desmedt and Elkind claim (Theorems 1 and 2) that for R^C and lazy voters a PNE exists if and only if the utility profile satisfies either condition (1) or condition (3) with constraint $k \leq n/2$ removed. To see why this is incorrect, consider a 2-voter election over the candidate set $C = \{x, y, z\}$, where voters' utility functions are consistent with preference orders $x \succ y \succ z$ and $x \succ z \succ y$, respectively. According to Desmedt and Elkind, the ballot vector (y, z) is a PNE of the corresponding game. This is obviously not true: each of the voters would prefer to change her vote to x. Note, however, that the two characterizations differ only when $m \geq n$, and in practice the number of voters usually exceeds the number of candidates.

Desmedt and Elkind show that checking condition (3) of Theorem 2 is NP-hard; in their proof $n > m$, and the proof does not depend on how the trivial ballot is handled. Further, their proof shows that checking whether a given candidate belongs to some such set X is also NP-hard. On the other hand, Theorem 2 shows that PNE with singleton winning sets only arise if some candidate is unanimously ranked first, and this condition is easy to check. We summarize these observations as follows.

Corollary 2. *For* $R \in \{R^C, R^V\}$*, the problems* $(\mathcal{L}, R)$-EXISTNE *and* $(\mathcal{L}, R)$-TIENE *are* NP*-complete, whereas* $(\mathcal{L}, R)$-SINGLENE *is in* P.

4 Truth-biased Voters

For truth-biased voters, our exposition follows the same pattern as for lazy voters: we present some general observations, followed by a quick summary of the results for lexicographic tie-breaking, and conclude by analyzing randomized tie-breaking. The following result is similar in spirit to Proposition 1.

Proposition 2. *For every* $R \in \{R^L, R^C, R^V\}$ *and every utility profile* $\mathbf{u}$*, if a ballot vector* $\mathbf{b}$ *is a PNE of* $(\mathcal{T}, R, \mathbf{u})$ *then for every voter* $i \in N$ *either* $b_i = a_i$*, or* $b_i \in W(\mathbf{b})$.

Proof. Consider a voter $i \in N$ such that $a_i \neq b_i$ and $b_i \notin W(\mathbf{b})$. Suppose $a_i \notin W(\mathbf{b})$. Then, if i changes her vote to a_i, the new winning set is either $W(\mathbf{b})$ or $W(\mathbf{b}) \cup \{a_i\}$.

In either case, i's utility increases at least by ε, a contradiction. Suppose now that $a_i \in W(\mathbf{b})$. This means that either $W(\mathbf{b}) = \{a_i\}$ or a_i is in a tie with other candidates under $\mathbf{b}$. Then, if i votes for a_i, the new winning set is just $\{a_i\}$, so i's utility increases by at least ε, a contradiction again. □

Lexicographic Tie-breaking Obraztsova et al. [17] characterize the PNE of the game $(\mathcal{T}, R^L, \mathbf{u})$. As their characterization is quite complex, we will not reproduce it here. However, for the purposes of comparison with the lazy voters model, we will use the following description of *truthful* equilibria given by Obraztsova et al.

Proposition 3 (Obraztsova et al., Theorem 1)**.** *Consider a utility profile* $\mathbf{u}$*, let* $\mathbf{a}$ *be the respective truthful ballot vector, and let* $j = \min\{r \mid c_r \in W(\mathbf{a})\}$*. Then* $\mathbf{a}$ *is a PNE of* $(\mathcal{T}, R^L, \mathbf{u})$ *if and only if neither of the following conditions holds:*

(1) $|W(\mathbf{a})| > 1$*, and there exists a candidate* $c_k \in W(\mathbf{a})$ *and a voter* i *such that* $a_i \neq c_k$ *and* $c_k \succ_i c_j$*.*

(2) $H(\mathbf{a}) \neq \emptyset$*, and there exists a candidate* $c_k \in H(\mathbf{a})$ *and a voter* i *such that* $a_i \neq c_k$*,* $c_k \succ_i c_j$*, and* $k < j$*.*

We will also state a crucial property of non-truthful PNE, identified by Obraztsova et al. For this, we first need the following definition.

Definition 1. *Consider a ballot vector* $\mathbf{b}$*, where candidate* c_j *is the winner under* R^L*. A candidate* $c_k \neq c_j$ *is called a* threshold candidate with respect to $\mathbf{b}$ *if either (1)* $k < j$ *and* $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_j, \mathbf{b}) - 1$ *or (2)* $k > j$ *and* $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_j, \mathbf{b})$*. We denote the set of threshold candidates with respect to* $\mathbf{b}$ *by* $T(\mathbf{b})$*.*

That is, a threshold candidate is someone who could win the election if he had one additional vote. A feature of all non-truthful PNE is that there must exist at least one threshold candidate. The intuition for this is that, since voters who are not pivotal prefer to vote truthfully, in any PNE that arises under strategic voting, the winner receives just enough votes so as to beat the required threshold (as set by the threshold candidate) and not any more.

Lemma 1 (Obraztsova et al., Lemma 2)**.** *Consider a utility profile* $\mathbf{u}$*, let* $\mathbf{a}$ *be the respective truthful ballot vector, and let* $\mathbf{b} \neq \mathbf{a}$ *be a non-truthful PNE of* $(\mathcal{T}, R^L, \mathbf{u})$*. Then* $T(\mathbf{b}) \neq \emptyset$*. Further,* $\mathrm{sc}(c_k, \mathbf{b}) = \mathrm{sc}(c_k, \mathbf{a})$ *for every* $c_k \in T(\mathbf{b})$*, i.e., all voters whose top choice is* c_k *vote for* c_k*.*

The existence of a threshold candidate is an important observation about the structure of non-truthful PNE, and we will use it repeatedly in the sequel. We note that the winner in $\mathbf{a}$ need not necessarily be a threshold candidate in a non-truthful PNE $\mathbf{b}$.

Obraztsova et al. show that, given a candidate $c_p \in C$ and a score s, it is computationally hard to decide whether the game $(\mathcal{T}, R^L, \mathbf{u})$ has a PNE $\mathbf{b}$ where c_p wins with a score of s. This problem may appear to be "harder" than $(\mathcal{T}, R^L)$-TIENE or $(\mathcal{T}, R^L)$-SINGLENE, as one needs to ensure that c_p obtains a specific score; on the other hand, it does not distinguish between c_p being the unique top-scorer or being tied with other candidates and winning due to tie-breaking. We now complement this hardness result by showing that all three problems we consider are NP-hard for $\mathcal{T}$ and R^L.

Theorem 3. $(\mathcal{T}, R^L)$-SINGLENE, $(\mathcal{T}, R^L)$-EXISTNE, *and* $(\mathcal{T}, R^L)$-TIENE *are* NP*-complete.*

The proof is by a reduction from MAXIMUM k-SUBSET INTERSECTION (MSI) (see the supplementary material). Surprisingly, the complexity of MSI was very recently posed as an open problem by Clifford and Popa [5]; subsequently, MSI was shown to be hard under Cook reductions by Xavier [21]. Here we first establish NP-hardness of MSI under Karp reductions, which may be of independent interest, and then show NP-hardness of our problems by constructing reductions from MSI.

Randomized Tie-breaking It turns out that for truth-biased voters the tie-breaking rules R^C and R^V induce identical behavior by the voters; unlike for lazy voters, this holds even if all voters rank the same candidate first.

For clarity, we present our characterization result for randomized tie-breaking in three parts. We start by considering PNE with winning sets of size at least 2; the analysis for this case turns out to be very similar to that for lazy voters.

Theorem 4. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C, $|C| = m$, *and let* $R \in \{R^C, R^V\}$. *The game* $G = (\mathcal{T}, R, \mathbf{u})$ *admits a PNE with a winning set of size at least* 2 *if and only if one of the following conditions holds:*

(1) each candidate is ranked first by at most one voter, and, moreover, $\frac{1}{n}\sum_{i \in N} u_\ell(a_i) \geq \max_{i \in N \setminus \{\ell\}} u_\ell(a_i)$ *for each* $\ell \in N$.

(2) there exists a set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ *with* $2 \leq k \leq \min(n/2, m)$ *and a partitioning of the voters into* k *groups* $N_1, \ldots, N_k$, *of size* n/k *each, such that for each* $j \in [k]$ *and each* $i \in N_j$, *we have* $c_{\ell_j} \succ_i c$ *for all* $c \in X \setminus \{c_{\ell_j}\}$, *and, moreover,* $\frac{1}{k}\sum_{c \in X} u_i(c) \geq \max_{c \in X \setminus \{c_{\ell_j}\}} u_i(c)$.

Further, if condition (1) holds, then G *has a PNE where each voter votes for her top candidate, and if condition (2) holds for some* X, *then* G *has a PNE where each voter votes for her favorite candidate in* X. *The game* G *has no other PNE.*

The case where the winning set is a singleton is surprisingly complicated. We will first characterize utility profiles that admit a truthful PNE with this property.

Theorem 5. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C, *let* $R \in \{R^C, R^V\}$, *and suppose that* $W(\mathbf{a}) = \{c_j\}$ *for some* $c_j \in C$. *Then* $\mathbf{a}$ *is a PNE of the game* $G = (\mathcal{T}, R, \mathbf{u})$ *if and only if for every* $i \in N$ *and every* $c_k \in H(\mathbf{a}) \setminus \{a_i\}$, *it holds that* $c_j \succ_i c_k$.

Finally, we consider elections that have non-truthful equilibria with singleton winning sets.

Theorem 6. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C, *let* $R \in \{R^C, R^V\}$, *and consider a ballot vector* $\mathbf{b}$ *with* $W(\mathbf{b}) = \{c_j\}$ *for some* $c_j \in C$ *and* $b_r \neq a_r$ *for some* $r \in N$. *Then* $\mathbf{b}$ *is a PNE of the game* $G = (\mathcal{T}, R, \mathbf{u})$ *if and only if all of the following conditions hold:*

(1) $b_i \in \{a_i, c_j\}$ *for all* $i \in N$;

(2) $H(\mathbf{b}) \neq \emptyset$*;*

(3) $c_j \succ_i c_k$ *for all* $i \in N$ *and all* $c_k \in H(\mathbf{b}) \setminus \{b_i\}$*;*

(4) for every candidate $c_\ell \in H'(\mathbf{b})$ *and each voter* $i \in N$ *with* $b_i = c_j$*,* i *prefers* c_j *to the lottery where a candidate is chosen from* $H(\mathbf{b}) \cup \{c_j, c_\ell\}$ *according to* R*.*

We now consider the complexity of EXISTNE, TIENE, and SINGLENE for truth-biased voters and randomized tie-breaking. The reader may observe that the characterization of PNE with ties in Theorem 6 is essentially identical to the one in Theorem 2. As a consequence, we immediately obtain that $(\mathcal{T}, R^C)$-TIENE and $(\mathcal{T}, R^V)$-TIENE are NP-hard. For EXISTNE and SINGLENE, a simple modification of the proof of Theorem 3 shows that these problems remain hard under randomized tie-breaking. These observations are summarized in the following corollary.

Corollary 3. *For* $R \in \{R^C, R^V\}$*,* $(\mathcal{T}, R)$-SINGLENE*,* $(\mathcal{T}, R)$-TIENE*, and* $(\mathcal{T}, R)$-EXISTNE *are* NP*-complete.*

5 Comparison

We are finally in a position to compare the different models considered in this paper.

Tie-breaking rules We have demonstrated that in equilibrium the two randomized tie-breaking rules (R^C and R^V) induce very similar voter behavior, and identical election outcomes, both for lazy and for truth-biased voters. This is quite remarkable, since under truthful voting these tie-breaking rules can result in very different lotteries. In contrast, there is a substantial difference between the randomized rules and the lexicographic rule. For instance, when voters are lazy, EXISTNE is NP-hard for R^C and R^V, but polynomial-time solvable for R^L. Further, the lexicographic rule is, by definition, not anonymous, and Theorem 1 demonstrates that candidates with smaller indices have a substantial advantage. For truth-biased voters the impact of tie-breaking rules is less clear: while we have obtained NP-hardness results for all three rules, it appears that, in contrast with lazy voters, for truth-biased voters randomized tie-breaking induces "simpler" PNE than lexicographic tie-breaking.

Lazy vs. truth-biased voters Under lexicographic tie-breaking, the sets of equilibria induced by the two types of secondary preferences are incomparable: there exists a utility profile $\mathbf{u}$ such that the sets of candidates who can win in PNE of $(\mathcal{L}, R^L, \mathbf{u})$ and $(\mathcal{T}, R^L, \mathbf{u})$ are disjoint.

Example 1. Let $C = \{c_1, c_2, c_3\}$, and consider a 4-voter election with one vote of the form $c_2 \succ c_3 \succ c_1$, and three votes of the form $c_3 \succ c_2 \succ c_1$. The only PNE of $(\mathcal{L}, R^L, \mathbf{u})$ is $(c_2, \perp, \perp, \perp)$, where c_2 wins, whereas the only PNE of $(\mathcal{T}, R^L, \mathbf{u})$ is (c_2, c_3, c_3, c_3), where c_3 wins.

For randomized tie-breaking, the situation is more interesting. For concreteness, let us focus on R^C. Note first that the utility profiles for which there exist PNE with winning sets of size 2 or more are the same for both voter types. Further, if $(\mathcal{L}, R^C, \mathbf{u})$ has a PNE $\mathbf{b}$ with $|W(\mathbf{b})| = 1$ (which happens only if there is a unanimous winner),

then $\mathbf{b}$ is also a PNE of $(\mathcal{T}, R^C, \mathbf{u})$. However, $(\mathcal{T}, R^C, \mathbf{u})$ may have additional PNE, including some non-truthful ones. In particular, for truth-biased voters, the presence of a strong candidate is sufficient for stability: Proposition 3 implies that if there exists a $c \in C$ such that $\mathrm{sc}(c, \mathbf{a}) \geq \mathrm{sc}(c', \mathbf{a}) + 2$ for all $c' \in C \setminus \{c\}$, then for any $R \in \{R^L, R^C, R^V\}$ the ballot vector $\mathbf{a}$ is a PNE of $(\mathcal{T}, R, \mathbf{u})$ with $W(\mathbf{a}) = \{c\}$.

Existence of PNE One can argue that, when the number of voters is large relative to the number of candidates, under reasonable probabilistic models of elections, the existence of a strong candidate (as defined in the previous paragraph) is exceedingly likely (we omit the formal statement of this result and its proof due to space constraints), so elections with truth-biased voters typically admit stable outcomes; this is corroborated by the experimental results of [20]. In contrast, for lazy voters stability is more difficult to achieve, unless there is a candidate that is unanimously ranked first: under randomized tie-breaking rules, there needs to be a very precise balance among the candidates that end up being in $W(\mathbf{b})$, and under R^L the eventual winner has to Pareto-dominate all candidates that lexicographically precede him. Either of these conditions appears to be quite difficult to satisfy in a large election.

Quality of PNE In all of our models, a candidate ranked last by all voters cannot be elected, in contrast to the basic game-theoretic model for Plurality voting. However, not all non-desirable outcomes are eliminated: under R^V and R^C both lazy voters and truth-biased voters can still elect a Pareto-dominated candidate with non-zero probability in PNE. This has been shown for lazy voters and R^C by [6] (Example 1), and the same example works for truth-biased voters and for R^V. A similar construction shows that a Pareto-dominated candidate may win under R^L when voters are truth-biased. In contrast, lazy voters cannot elect a Pareto-dominated candidate under R^L: Theorem 1 shows that the winner has to be ranked first by some voter.

We can also measure the quality of PNE by analyzing the Price of Anarchy (PoA) in both models. The study of PoA in the context of voting has been recently initiated by Branzei et al. [4]. The additive version of PoA, which was considered by Branzei et al., is defined as the worst-case difference between the score of the winner under truthful voting and the truthful score of a PNE winner. It turns out that PoA can be quite high, both for lazy and truth-biased voters. To illustrate this, we provide in the supplementary material two examples showing that under lexicographic tie-breaking $\mathrm{PoA} = \Omega(n)$ in both models. Similar results can be established for randomized tie-breaking as well.

Even though the PoA results are not encouraging, this is only a worst-case analysis and we expect PNE to have a better performance on average. For the truth-biased model, this is also supported by the experimental evaluation of Thompson et al. [20], who showed that in the truth-biased model most PNE identified in their simulations had good social welfare properties. Formalizing this observation, i.e., providing average-case analysis of the quality of PNE in voting games, is a promising topic for future work.

6 Extension: Principled Voters

The results of this paper can be extended to the setting where some of the voters are *principled*, i.e., always vote truthfully (and never abstain). Due to space constraints,

we relegate the formal statements of our results for this extended model to the supplementary material. Briefly, the presence of principled voters has the strongest effect on lazy voters and lexicographic tie-breaking, whereas for other settings the effect is less pronounced. All computational problems that were easy in the standard model remain easy in the extended model (and, obviously, all hard problems remain hard). Finally, in the presence of principled voters the random candidate tie-breaking rule is no longer equivalent to the random voter tie-breaking rule.

7 Conclusions and Future Work

We have characterized PNE of Plurality voting for several combinations of secondary preferences and tie-breaking rules. Our complexity results are summarized in Table 1. A promising direction for future work is to investigate more general classes of tie-breaking rules. It is also interesting to consider the complexity of various refinements of Nash equilibria for our models, such as strong Nash equilibria (for which an analysis for $\mathcal{T}$ and R^L can be found in the work of Obraztsova et al. [17]), or subgame-perfect Nash equilibria for settings where voters submit their ballots one by one; see [6] and [22] for some results about such equilibria.

	SINGLENE	TIENE	EXISTNE
$(\mathcal{L}, R^L)$	P (Cor. 1)	P (Cor. 1)	P (Cor. 1)
$(\mathcal{L}, R^C)$	P (Cor. 2)	NPc (Cor. 2)	NPc (Cor. 2)
$(\mathcal{L}, R^V)$	P (Cor. 2)	NPc (Cor. 2)	NPc (Cor. 2)
$(\mathcal{T}, R^L)$	NPc (Thm. 3)	NPc (Thm. 3)	NPc (Thm. 3)
$(\mathcal{T}, R^C)$	NPc (Cor. 3)	NPc (Cor. 3)	NPc (Cor. 3)
$(\mathcal{T}, R^V)$	NPc (Cor. 3)	NPc (Cor. 3)	NPc (Cor. 3)

Table 1: Complexity results: P stands for “polynomial-time solvable”, NPc stands for “NP-complete”.

References

[1] H. Aziz, S. Gaspers, N. Mattei, N. Narodytska, and T. Walsh. Ties matter: Complexity of manipulation when tie-breaking with a random vote. In *AAAI'13*, pages 74–80, 2013.

[2] M. Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51:445–463, 2005.

[3] T. Borgers. Costly voting. *American Economic Review*, 94(1):57–66, 2004.

[4] S. Branzei, I. Caragiannis, J. Morgenstern, and A. D. Procaccia. How bad is selfish voting? In *AAAI*, pages 138–144, 2013.

[5] R. Clifford and A. Popa. Maximum subset intersection. *Information Processing Letters*, 111(7):323–325, 2011.

[6] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC'10*, pages 347–356, 2010.

[7] A. Dhillon and B. Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.

[8] B. Dutta and A. Sen. Nash implementation with partially honest individuals. *Games and Economic Behavior*, 74(1):154–169, 2012.

[9] R. Farquharson. *Theory of Voting*. Yale University Press, 1969.

[10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1979.

[11] M. Lombardi and N. Yoshihara. A full characterization of Nash implementation with strategy space reduction. *Economic Theory*, 54(1):131–151, 2013.

[12] R. Meir, M. Polukarov, J. S. Rosenschein, and N. R. Jennings. Convergence to equilibria of plurality voting. In *AAAI'10*, pages 823–828, 2010.

[13] H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.

[14] R. Myerson and R. Weber. A theory of voting equilibria. *American Political Science Review*, 87(1):102–114, 1993.

[15] S. Obraztsova and E. Elkind. On the complexity of voting manipulation under randomized tie-breaking. In *IJCAI'11*, pages 319–324, 2011.

[16] S. Obraztsova, E. Elkind, and N. Hazon. Ties matter: Complexity of voting manipulation revisited. In *AAMAS'11*, pages 71–78, 2011.

[17] S. Obraztsova, E. Markakis, and D. R. M. Thompson. Plurality voting with truth-biased agents. In *SAGT'13*, pages 26–37, 2013.

[18] M. Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2):267–303, 2011.

[19] F. D. Sinopoli and G. Iannantuoni. On the generic strategic stability of Nash equilibria if voting is costly. *Economic Theory*, 25(2):477–486, 2005.

[20] D. R. M. Thompson, O. Lev, K. Leyton-Brown, and J. S. Rosenschein. Empirical analysis of plurality election equilibria. In *AAMAS'13*, pages 391–398, 2013.

[21] E. Xavier. A note on a maximum k-subset intersection problem. *Information Processing Letters*, 112(12):471–472, 2012.

[22] L. Xia and V. Conitzer. Stackelberg voting games: Computational aspects and paradoxes. In *AAAI'10*, pages 805–810, 2010.

A Proofs Omitted from Section 3

Theorem 1. *For any utility profile* $\mathbf{u}$ *the game* $G = (\mathcal{L}, R^L, \mathbf{u})$ *has the following properties:*

1. *If* $\mathbf{b}$ *is a PNE of* G *then* $|W(\mathbf{b})| \in \{1, m\}$*. Moreover,* $|W(\mathbf{b})| = m$ *if and only if* $\mathbf{b}$ *is the trivial ballot and all voters rank* c_1 *first.*
2. *If* $\mathbf{b}$ *is a PNE of* G *then there exists at most one voter* i *with* $b_i \neq \bot$*.*
3. G *admits a PNE if and only if all voters rank* c_1 *first (in which case* c_1 *is the unique PNE winner) or there exists a candidate* c_j *with* $j > 1$ *such that (i)* $\mathrm{sc}(c_j, \mathbf{a}) > 0$ *and (ii) for every* $k < j$ *it holds that all voters prefer* c_j *to* c_k*. If such a candidate exists, he is unique, and wins in all PNE of* G*.*

Proof. Fix a utility profile $\mathbf{u}$ and a ballot $\mathbf{b}$ such that $\mathbf{b}$ is a PNE of $G = (\mathcal{L}, R^L, \mathbf{u})$.

To prove the first claim, suppose first that $1 < |W(\mathbf{b})|$ and $\mathbf{b}$ is not trivial. Then there are two candidates $c_j, c_k \in W(\mathbf{b})$, $j < k$, such that $\mathrm{sc}(c_j, \mathbf{b}) > 0$ and $\mathrm{sc}(c_k, \mathbf{b}) > 0$. Hence, there exists at least one voter who votes for c_k. However, the election outcome will not change if this voter abstains, a contradiction with $\mathbf{b}$ being a PNE of G. Now, suppose that $\mathbf{b}$ is trivial. In this case $W(\mathbf{b}) = C$ and c_1 wins. If any voter prefers some other candidate c to c_1, she can improve her utility by voting for c, as this will change the election outcome to c. On the other hand, if all voters rank c_1 first, the trivial ballot is clearly a PNE.

The second claim follows from our first claim and Proposition 1.

To prove the third claim, suppose that there exists a candidate c_j, $j > 1$, satisfying conditions (i) and (ii). Consider a ballot vector $\mathbf{b}$ where $b_i = c_j$ for some voter i with $a_i = c_j$ (the existence of such voter is guaranteed by condition (i)) and $b_{i'} = \bot$ for all $i' \in N \setminus \{i\}$. Voter i cannot benefit from voting for another candidate or abstaining, as this will change the election outcome to one she likes less than the current outcome. Any other voter can only change the election outcome if she votes for a candidate c_k with $k < j$. But then condition (ii) implies that no voter wants the election outcome to change in this way. Conversely, suppose that $\mathbf{b}$ is a PNE. We have argued that either $\mathbf{b}$ is trivial or $b_i = c_j$ for some $i \in N$ and some $c_j \in C$ and $b_{i'} = \bot$ for all $i' \in N \setminus \{i\}$. In the latter case, if $c_j \neq a_i$, voter i can improve her utility by voting for a_i. Moreover, if $j = 1$, voter i can improve her utility by abstaining, as c_1 would remain the election winner in this case. Finally, if there exists a candidate c_k with $k < j$ such that some voter i' prefers c_k to c_j, then i' can change the election outcome to c_k by voting for c_k.

It remains to show that conditions (i) and (ii) can be satisfied by at most one candidate. To see this, note that if both c_j and c_k satisfy condition (i) and $j < k$, then c_k violates condition (ii), as the voter who ranks c_j first clearly prefers c_j to c_k. □

Theorem 2. *Let* $\mathbf{u} = (u_1, \ldots, u_n)$ *be a utility profile over* C*,* $|C| = m$*, and let* $R \in \{R^C, R^V\}$*. The game* $G = (\mathcal{L}, R, \mathbf{u})$ *admits a PNE if and only if one of the following conditions holds:*

(1) all voters rank some candidate c_j *first;*

(2) each candidate is ranked first by at most one voter, and, moreover, $\frac{1}{n}\sum_{i\in N} u_\ell(a_i) \geq \max_{i\in N\setminus\{\ell\}} u_\ell(a_i)$ *for each* $\ell \in N$.

(3) there exists a set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ *with* $2 \leq k \leq \min(n/2, m)$ *and a partition of the voters into* k *groups* $N_1, \ldots, N_k$ *of size* n/k *each such that for each* $j \in [k]$ *and each* $i \in N_j$ *we have* $c_{\ell_j} \succ_i c$ *for all* $c \in X \setminus \{c_{\ell_j}\}$*, and, moreover,* $\frac{1}{k}\sum_{c\in X} u_i(c) \geq \max_{c\in X\setminus\{c_{\ell_j}\}} u_i(c)$.

Further, if condition (1) holds for some $c_j \in C$*, then if* $R = R^C$ *then for each* $i \in N$ *the game* G *has a PNE where* i *votes for* c_j *and all other voters abstain, whereas if* $R = R^V$ *the game* G *has a PNE where all voters abstain; if condition (2) holds, then* G *has a PNE where each voter votes for her top candidate; and if condition (3) holds for some set* X*, then* G *has a PNE where each voter votes for her favorite candidate in* X*. The game* G *has no other PNE.*

Proof. It is easy to see that any of the conditions (1)–(3) is sufficient for the existence of PNE, with ballot vectors described in the statement of the theorem witnessing this. We will now show that satisfying at least one of these conditions is necessary for the existence of a PNE, and that no other ballot vector is a PNE. Fix a tie-breaking rule $R \in \{R^C, R^V\}$, a utility profile $\mathbf{u}$, and suppose that a ballot vector $\mathbf{b}$ is a PNE of $(\mathcal{L}, R, \mathbf{u})$. We will argue that $\mathbf{u}$ satisfies one of the conditions (1)–(3).

Suppose first that $W(\mathbf{b}) = \{c_j\}$ for some $c_j \in C$. By Proposition 1 there exists a voter $i \in N$ with $b_i = c_j$, and $b_{i'} = \perp$ for all $i' \in N \setminus \{i\}$. It has to be the case that $a_i = c_j$: otherwise voter i can make a_i the unique winner by changing her vote to a_i, thus increasing her utility. Now, suppose that $a_{i'} \neq c_j$ for some $i' \in N \setminus \{i\}$. If voter i' changes her ballot to $c_\ell = a_{i'}$, the new winning set is $\{c_j, c_\ell\}$. Now, if $R = R^C$, the overall utility of i' is given by $\frac{1}{2}(u_{i'}(c_\ell) + u_{i'}(c_j))$, and if $R = R^V$, the overall utility of i' is given by $\lambda u_{i'}(c_\ell) + (1-\lambda)u_{i'}(c_j)$, where $\lambda \geq \frac{1}{n}$ (this is because voter i' herself ranks c_ℓ above c_j). In both cases, i' can increase her utility by voting c_ℓ, a contradiction. Hence, it has to be the case that all voters rank c_j first, i.e., condition (1) is satisfied.

Now, suppose that $|W(\mathbf{b})| > 1$. We will argue that in this case either all voters abstain or no voter abstains. Indeed, suppose that $b_i = \perp$, $b_\ell \neq \perp$ for some $i, \ell \in N$, i.e., each candidate in $W(\mathbf{b})$ receives at least one vote. If, instead of abstaining, i votes for her most preferred candidate in $W(\mathbf{b})$, this candidate becomes the unique election winner. In contrast, under $\mathbf{b}$ i's least preferred candidate in $W(\mathbf{b})$ wins with positive probability: this is immediate for $R = R^C$, and for $R = R^V$ this holds because for every $c_j \in W(\mathbf{b})$ there exists a voter i' with $b_{i'} = c_j$, and c_j wins whenever ties are broken according to the preferences of voter i'. Thus, i can improve her utility by changing her vote, a contradiction. Hence, if $|W(\mathbf{b})| = k$ and $\mathbf{b}$ is not trivial, each candidate in $W(\mathbf{b})$ receives exactly n/k votes.

In particular, if $|W(\mathbf{b})| = n$ and $\mathbf{b}$ is not trivial, each candidate in $W(\mathbf{b})$ receives exactly one vote. We will argue that in this case condition (2) is satisfied. We will first prove that $b_i = a_i$ for all $i \in N$. Indeed, suppose that $b_i \neq a_i$ for some $i \in N$, and consider the ballot vector $\mathbf{b}' = (\mathbf{b}_{-i}, a_i)$. If $a_i \in W(\mathbf{b})$, then $W(\mathbf{b}') = \{a_i\}$, whereas under $\mathbf{b}$ voter i's least preferred candidate in $W(\mathbf{b})$ wins with positive probability. If $a_i \notin W(\mathbf{b})$, we have $W(\mathbf{b}') = (W(\mathbf{b}) \setminus \{b_i\}) \cup \{a_i\}$, so $U_i(\mathbf{b}') = U_i(\mathbf{b}) + \frac{1}{n}(u_i(a_i) -$

$u_i(b_i)) > U_i(\mathbf{b})$. In both cases i can increase her overall utility by voting for a_i, a contradiction. Hence, we have $W(\mathbf{b}) = \{a_i \mid i \in N\}$. Thus, under both R^C and R^V the outcome of this election is a lottery that assigns equal probability to all candidates in $W(\mathbf{b})$. Now, if any voter prefers her second most preferred candidate in $W(\mathbf{b})$ to this lottery, she can vote for that candidate, making him the unique election winner, a contradiction with $\mathbf{b}$ being a PNE. Thus, in this case condition (2) is satisfied.

Now, suppose that $\mathbf{b}$ is not trivial and $|W(\mathbf{b})| = k < n$. We have argued that each candidate in $W(\mathbf{b})$ receives exactly n/k votes. This means that k divides n, so in particular $k \leq n/2$ and each candidate in $W(\mathbf{b})$ receives at least two votes. Under both of our tie-breaking rules, each candidate in $W(\mathbf{b})$ wins with probability $1/k$. Consider a voter i. She can make any candidate in $W(\mathbf{b}) \setminus \{b_i\}$ the unique election winner by voting for him. Since $\mathbf{b}$ is a PNE, no voter wants to change the election outcome in this way; this implies, in particular, that each voter votes for her favorite candidate in $W(\mathbf{b})$. Thus, in this case condition (3) is satisfied with $X = W(\mathbf{b})$; the voters are partitioned into groups according to their votes in $\mathbf{b}$.

It remains to consider the case where $\mathbf{b}$ is the trivial ballot vector. When $R = R^C$, $\mathbf{b}$ cannot be a PNE: under $\mathbf{b}$ the outcome is a uniform lottery over C, and every voter would rather vote for her favorite candidate in order to make him the unique winner. When $R = R^V$, the outcome is a lottery that assigns a positive probability to each candidate in $A = \{a_i \mid i \in N\}$. If $|A| > 1$, $\mathbf{b}$ is not a PNE: each voter would prefer to vote for her favorite candidate in order to make him the unique winner. However, if A is a singleton, i.e., all voters rank some candidate c_j first, the trivial ballot vector is a PNE: after all voters abstain, R^V picks a random voter, and this voter selects c_j. □

B Proofs Omitted from Section 4

The following problem is very useful in our constructions.

Definition 2. *An instance of* MAXIMUM k-SUBSET INTERSECTION (MSI) *is a tuple* $(\mathcal{E}, A_1, \ldots, A_m, k, q)$*, where* $\mathcal{E} = \{e_1, \ldots, e_n\}$ *is a finite set of elements, each* A_i*,* $i \in [m]$*, is a subset of* $\mathcal{E}$*, and* k, q *are positive integers. It is a "yes"-instance if there exist sets* $A_{i_1}, \ldots, A_{i_k}$ *such that* $|\cap_{j\in[k]} A_{i_j}| \geq q$*, and a "no"-instance otherwise.*

Despite the relevance of MSI to various optimization scenarios, see e.g. [21], it was only recently shown that this problem is hard under a Cook reduction. We provide below a Karp reduction, establishing NP-completeness of MSI.

Theorem 7. MSI *is* NP*-complete.*

Proof. Trivially MSI is in NP. For hardness, we provide a reduction from the BALANCED COMPLETE BIPARTITE SUBGRAPH problem. An instance of this problem consists of a bipartite graph $G = (V_1, V_2, E)$, and an integer k. It is a "yes"-instance if there exist sets $S_1 \subseteq V_1$, $S_2 \subseteq V_2$, with $|S_1| = |S_2| = k$, such that the subgraph induced by S_1 and S_2 is a complete bipartite subgraph. This problem is known to be NP-complete, see [10].

Consider an instance I of this problem, given by G and some integer k. We construct an instance I' of MSI as follows: we let $\mathcal{E} = V_1$. We also let the elements of V_2 correspond to sets. In particular, for every $j \in V_2$, we have a corresponding set $A_j \subseteq \mathcal{E}$, such that $A_j = \{i \in V_1 : (i,j) \in E\}$. We set the parameter k in MSI to be the same as the integer k from I. We also set $q = k$. We now claim that I is a "yes"-instance of BALANCED COMPLETE BIPARTITE SUBGRAPH if and only if I' is a "yes"-instance of MSI.

To see this, suppose there exist $S_1 \subseteq V_1$, $S_2 \subseteq V_2$, such that we have a complete bipartite subgraph induced by S_1 and S_2. Then take the k sets corresponding to S_2. Clearly every element from S_1 belongs to all these sets, hence the intersection of these sets is at least k. For the reverse direction, suppose there exists a collection of k sets in I' whose intersection is at least k. Then take as S_2 the k vertices that correspond to these sets. Let also S_1 be any k-element subset of the intersection. Then obviously, the bipartite graph induced by S_1 and S_2 is complete. □

We can now prove Theorem 3, utilizing the hardness of MSI.

Theorem 3. $(\mathcal{T}, R^L)$-SINGLENE, $(\mathcal{T}, R^L)$-EXISTNE, *and* $(\mathcal{T}, R^L)$-TIENE *are* NP*-complete.*

Proof. We will first establish the NP-completeness of $(\mathcal{T}, R^L)$-SINGLENE, and then show how to modify the proof for the other two problems. It is trivial to show that $(\mathcal{T}, R^L)$-SINGLENE is in NP, so we focus on showing that it is NP-hard.

We provide a reduction from MSI. Consider an instance I of the MSI problem. We can assume that for every $e \in \mathcal{E}$ there exists an index i, such that $e \notin A^i$ and $m > n + k + q$ (if this is not the case we can add several additional empty sets).

We now construct an instance of our problem. We have $n + 3$ candidates, namely, $C = \mathcal{E} \cup \{w_1, w_2, w_3\}$, with ties broken according to $e_1 > \ldots e_n > w_3 > w_1 > w_2$. We set w_2 to be the target winning candidate, i.e., $c_p := w_2$. Finally, we set $\delta = \frac{1}{6(n+m)}$. We will now describe the voters' preferences and their utility functions (while the utility functions play no role in this proof, we will use the same construction in the NP-hardness proof for randomized tie-breaking, see Corollary 3, where they do matter). The voters in our instance are split into five blocks as follows.

- **Block 1:** For every $i \in [m]$ we construct a voter v_i who ranks the candidates as $w_3 \succ \mathcal{E} \setminus A^i \succ w_2 \succ A^i \succ w_1$. We let u_i denote the utility function of v_i, which we construct as follows. We set $u_i(w_3) = 1, u_i(w_2) = \frac{1}{2}, u_i(w_1) = \frac{1}{4}$. Further, v_i assigns utility of $1 - j\delta$ to her j-th most preferred candidate in $\mathcal{E} \setminus A^i$. Note that $|\mathcal{E} \setminus A^i| < n$, so these numbers are strictly between 1 and $1/2$ and they are consistent with the ranking of voter v_i. Finally, v_i assigns utility of $1/2 - j\delta$ to her j-th most preferred candidate in A^i; these numbers are strictly between $1/2$ and $1/4$.

- **Block 2:** We set $s = m - k + 3$, and we add $s - 1$ voters whose preferences are of the form $w_1 \succ w_2 \succ w_3 \succ \mathcal{E}$.

- **Block 3:** We add $s - k - (n - q) - 1$ voters with preferences of the form $w_2 \succ w_1 \succ w_3 \succ \mathcal{E}$.

- **Block 4:** For every $e_j \in \mathcal{E}$, we add $s-2$ voters with preferences of the form $e_j \succ w_3 \succ w_2 \succ w_1 \succ \mathcal{E} \setminus \{e_j\}$.
- **Block 5:** A voter with preferences of the form $w_3 \succ w_2 \succ \mathcal{E} \succ w_1$.

Each of the voters in blocks 2–5 assigns utility of $1-\delta(j-1)$ to the j-th candidate in her ranking.

Let I' be the constructed instance. We want to establish that I is a "yes"-instance of the MSI problem if and only if I' is a "yes"-instance of $(\mathcal{T}, R^L)$-SINGLENE. Suppose first that I' is a "yes"-instance of our problem. Then there exists a PNE $\mathbf{b}$ with $W(\mathbf{b}) = \{w_2\}$. We will first establish some properties of $\mathbf{b}$.

Let $\mathbf{a}$ denote the truthful ballot for I'. We have $\mathrm{sc}(w_1, \mathbf{a}) = s-1$, $\mathrm{sc}(w_2, \mathbf{a}) = s-k-(n-q)-1$, $\mathrm{sc}(w_3, \mathbf{a}) = m+1 = s+k-2$, and $\mathrm{sc}(e_j, \mathbf{a}) = s-2$ for every $e_j \in \mathcal{E}$. It follows that w_2 is not among the winners in $\mathbf{a}$.

We will now argue that in the PNE $\mathbf{b}$, $T(\mathbf{b}) = \{w_1\}$. We know by Lemma 1 that $T(\mathbf{b}) \neq \emptyset$. Since w_2 is the winner at $\mathbf{b}$, $w_2 \notin T(\mathbf{b})$. Also, it is easy to see that $w_3 \notin T(\mathbf{b})$. Indeed, suppose the contrary. All voters in Block 4 prefer w_3 to w_2. By Proposition 2, in $\mathbf{b}$ these voters vote either for their top choice or for $W(\mathbf{b}) = \{w_2\}$, hence not for w_3. But if $w_3 \in T(\mathbf{b})$, each of these voters would prefer to switch to voting w_3, a contradiction with $\mathbf{b}$ being a PNE. A similar argument shows that $\mathcal{E} \cap T(\mathbf{b}) = \emptyset$. Indeed, we assumed that for every $e_\ell \in \mathcal{E}$ there exists i such that $e_\ell \notin A^i$. Then voter v_i from Block 1 prefers e_ℓ to w_2, and e_ℓ is not her top choice. By Proposition 2, in $\mathbf{b}$ voter v_i votes for her top choice or for w_2, but if $e_\ell \in T(\mathbf{b})$, she would prefer to change her vote to e_ℓ, a contradiction with $\mathbf{b}$ being a PNE. As we have ruled out all candidates except for w_1, it follows that $T(\mathbf{b}) = \{w_1\}$, and hence, by Lemma 1, $\mathrm{sc}(w_1, \mathbf{b}) = s-1$. Then, by the tie-breaking rule, it must be that $\mathrm{sc}(w_2, \mathbf{b}) = s$. Thus, in $\mathbf{b}$ candidate w_2 receives exactly $k+n-q+1$ non-truthful votes, in addition to the votes of his own supporters. We also know that the voters from Block 3 keep voting for w_2 in $\mathbf{b}$, and, by Lemma 1, the voters from Block 2 keep voting for w_1 in $\mathbf{b}$. Hence w_2 receives the extra $k+n-q+1$ votes in $\mathbf{b}$ from Blocks 1, 4 and 5.

We claim that $\mathrm{sc}(w_3, \mathbf{b}) \leq s-3$. Indeed, we have $\mathrm{sc}(c', \mathbf{b}) \leq s-2$ for all $c' \in \mathcal{E} \cup \{w_3\}$ since $T(\mathbf{b}) = \{w_1\}$. Further, if $\mathrm{sc}(w_3, \mathbf{b}) = s-2$, then $\mathbf{b}$ would not be a PNE, as some voters from Blocks 1, 4, and 5 vote for w_2, but all of them prefer w_3 to w_2. Thus, since the only supporters of w_3 are in Block 1 and Block 5, in total, we must have at least $k+1$ voters from Blocks 1 and 5 who vote for w_2 in $\mathbf{b}$. This means that there are at least k voters from Block 1, who have deviated to w_2. Now, we pick all sets A_{i_j} for all $j \in [k]$ such that v_{i_j} votes for w_2 in $\mathbf{b}$. We will argue that $|\cap_{j\in[k]} A_{i_j}| \geq q$.

To see this, let $\mathcal{E}' = \{e \in \mathcal{E} \mid \mathrm{sc}(e, \mathbf{b}) = s-2\}$. Note that in $\mathbf{b}$ there are at most $n-q$ voters in Block 4 who vote for w_2. Hence, we have $|\mathcal{E}'| \geq q$. To complete the proof, we only need to argue that for each $e \in \mathcal{E}'$ we have $e \in A_{i_j}$ for all $j \in [k]$. Indeed, fix some $e \in \mathcal{E}'$ and some $j \in [k]$. By our choice of A_{i_j}, the corresponding voter v_{i_j} has voted for w_2 in $\mathbf{b}$. Suppose that v_{i_j} prefers e to w_2. If she changes her vote to e, then e becomes the new winner, due to tie-breaking, a contradiction with $\mathbf{b}$ being a PNE. Thus, it has to be the case that v_{i_j} prefers w_2 to e, which means that $e \in A_{i_j}$, as we wanted to prove. Hence, a "yes"-instance for $(\mathcal{T}, R^L)$-SINGLENE, corresponds to a "yes"-instance of MSI.

For the converse direction, suppose that for the instance I, there is a collection of sets $A_{i_1}, \ldots, A_{i_k}$ with $|\cap_{j\in[k]} A_{i_j}| \geq q$. Let $\mathcal{E}' = \cap_{j\in[k]} A_{i_j}$. We identify below a set of voters, N', from the instance I', which allows us to construct an equilibrium profile. We include in N' the set $\{v_{i_j} \mid j \in [k]\}$. We also add to N' the voter from Block 5. Furthermore, for each $e \notin \mathcal{E}'$, we add to N' one voter from Block 4 who ranks e first. Observe that at this point we have $|N'| \leq k+(n-q)+1$. If $|N'| < k+(n-q)+1$, we pick $n-q+k+1-|N'|$ additional voters from Block 4, corresponding to supporters of elements $e \notin \mathcal{E}'$, and add them to N'. Now, consider a ballot vector $\mathbf{b}$ where the voters in N' vote in favor of w_2, and everyone else votes truthfully. We have $\mathrm{sc}(w_2, \mathbf{b}) = s$, $\mathrm{sc}(w_1, \mathbf{b}) = s-1$, $\mathrm{sc}(w_3, \mathbf{b}) = s-3$, $\mathrm{sc}(e, \mathbf{b}) \leq s-3$ for all $e \in \mathcal{E}\setminus\mathcal{E}'$, $\mathrm{sc}(e, \mathbf{b}) = s-2$ for all $e \in \mathcal{E}'$. Hence, w_2 is the winner, and all non-truthful voters rank w_2 above w_1 as well as above all candidates in $\mathcal{E}'$ (who could possibly become winners if some voters had an incentive to vote for them). Thus, $\mathbf{b}$ is a PNE.

Finally, we comment on the hardness of the problems $(\mathcal{T}, R^L)$-EXISTNE and $(\mathcal{T}, R^L)$-TIENE. For $(\mathcal{T}, R^L)$-TIENE, we can make a small modification to the reduction above. Specifically, it suffices to switch the tie-breaking order between w_1 and w_2, and also add one more voter to Block 2 in favor of w_1. For $(\mathcal{T}, R^L)$-EXISTNE, hardness is again based on a modification of the reduction above: we can add $k-1$ additional copies of w_3 into the profile and change s to $m+2$. In this case it can be shown that w_1 is the only possible threshold candidate and hence only w_2 can be a winner in a PNE. We omit further details from this version. □

Theorem 4. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be a utility profile over C, $|C| = m$, and let $R \in \{R^C, R^V\}$. The game $G = (\mathcal{T}, R, \mathbf{u})$ admits a PNE with a winning set of size at least 2 if and only if one of the following conditions holds:*

(1) each candidate is ranked first by at most one voter, and, moreover, $\frac{1}{n}\sum_{i\in N} u_\ell(a_i) \geq \max_{i\in N\setminus\{\ell\}} u_\ell(a_i)$ for each $\ell \in N$.

(2) there exists a set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ with $2 \leq k \leq \min(n/2, m)$ and a partitioning of the voters into k groups $N_1, \ldots, N_k$, of size n/k each, such that for each $j \in [k]$ and each $i \in N_j$, we have $c_{\ell_j} \succ_i c$ for all $c \in X \setminus \{c_{\ell_j}\}$, and, moreover, $\frac{1}{k}\sum_{c\in X} u_i(c) \geq \max_{c\in X\setminus\{c_{\ell_j}\}} u_i(c)$.

Further, if condition (1) holds, then G has a PNE where each voter votes for her top candidate, and if condition (2) holds for some X, then G has a PNE where each voter votes for her favorite candidate in X. The game G has no other PNE.

Proof. It is clear that if one of the conditions (1)–(2) is satisfied then the game admits a PNE of the form described in the statement of the theorem. For the converse direction, fix a tie-breaking rule $R \in \{R^C, R^V\}$ and a utility profile $\mathbf{u}$, and suppose that a ballot vector $\mathbf{b}$ is a PNE of $(\mathcal{T}, R, \mathbf{u})$ with $|W(\mathbf{b})| \geq 2$. We will argue that $\mathbf{u}$ satisfies one of the conditions (1)–(2).

If $|W(\mathbf{b})| = n$, each candidate in $W(\mathbf{b})$ receives exactly one vote. As argued in the proof of Theorem 2, this means that each voter votes for her favorite candidate, and prefers the uniform lottery over $A = \{a_i \mid i \in N\}$ to her second most preferred candidate in A being the unique winner, i.e., condition (1) holds.

Now, suppose that $|W(\mathbf{b})| < n$. We claim that $b_i \in W(\mathbf{b})$ for all $i \in N$. Indeed, suppose that $b_i \notin W(\mathbf{b})$ for some $i \in N$. Let c_j be voter i's most preferred candidate in $W(\mathbf{b})$. If i changes her vote to c_j, c_j becomes the unique winner, whereas when she votes b_i, the outcome is a lottery over $W(\mathbf{b})$ where candidates other than c_j have a positive chance of winning. Thus, i can profitably deviate, a contradiction. Thus, there exists a $k \geq 2$ such that each candidate in $W(\mathbf{b})$ receives n/k votes. An argument similar to the one in the proof of Theorem 2 shows that condition (2) must be satisfied. □

Theorem 5. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be a utility profile over C, let $R \in \{R^C, R^V\}$, and suppose that $W(\mathbf{a}) = \{c_j\}$ for some $c_j \in C$. Then $\mathbf{a}$ is a PNE of the game $G = (\mathcal{T}, R, \mathbf{u})$ if and only if for every $i \in N$ and every $c_k \in H(\mathbf{a}) \setminus \{a_i\}$ it holds that $c_j \succ_i c_k$.*

Proof. Consider the ballot vector $\mathbf{a}$ and a voter $i \in N$. Clearly, if $a_i = c_j$, voter i cannot improve her utility by deviating. Otherwise, the only way i can change the election outcome is by changing her vote to some $c_k \in H(\mathbf{a}) \setminus \{a_i\}$, in which case the outcome is a lottery over $\{c_j, c_k\}$ where both of these candidates has a positive chance of winning. The condition of the theorem says that no voter wants to change the election outcome in this way. □

Theorem 6. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be a utility profile over C, let $R \in \{R^C, R^V\}$, and consider a ballot vector $\mathbf{b}$ with $W(\mathbf{b}) = \{c_j\}$ for some $c_j \in C$ and $b_r \neq a_r$ for some $r \in N$. Then $\mathbf{b}$ is a PNE of the game $G = (\mathcal{T}, R, \mathbf{u})$ if and only if all of the following conditions hold:*

(1) $b_i \in \{a_i, c_j\}$ for all $i \in N$;

(2) $H(\mathbf{b}) \neq \emptyset$;

(3) $c_j \succ_i c_k$ for all $i \in N$ and all $c_k \in H(\mathbf{b}) \setminus \{b_i\}$;

(4) for every candidate $c_\ell \in H'(\mathbf{b})$ and each voter $i \in N$ with $b_i = c_j$, i prefers c_j to the lottery where a candidate is chosen from $H(\mathbf{b}) \cup \{c_j, c_\ell\}$ according to R.

Proof. Suppose that a ballot profile $\mathbf{b}$ satisfies conditions (1)–(4) of the theorem, and consider a voter $i \in N$. If $b_i = a_i = c_j$, the current outcome is optimal for i. If $b_i = a_i \neq c_j$, the only way that voter i can change the election outcome is by voting for a candidate $c_k \in H(\mathbf{b}) \setminus \{a_i\}$, in which case the winner will be chosen from $\{c_j, c_k\}$ according to R. By condition (3), voter i does not benefit from this change. By Proposition 2, the only remaining possibility is that $b_i = c_j \neq a_i$. Then i can change the election outcome by (a) voting for a candidate $c_k \in H(\mathbf{b})$; (b) voting for a candidate $c_\ell \in H'(\mathbf{b})$; or (c) voting for a candidate in $C \setminus (H(\mathbf{b}) \cup H'(\mathbf{b}) \cup \{c_j\})$. In case (a) c_k becomes the unique winner, so by condition (3) this change is not profitable to i. In case (b) the outcome is a tie among the candidates in $H(\mathbf{b}) \cup \{c_j, c_\ell\}$, so by condition (4) voter i cannot profit from this change. Finally, in case (c) the outcome is a tie among the candidates in $H(\mathbf{b}) \cup \{c_j\}$, and by condition (3), i prefers the current outcome to this one. Thus, a ballot vector satisfying conditions (1)–(4) is indeed a PNE.

Conversely, suppose that $\mathbf{b}$ is a PNE of $(\mathcal{T}, R, \mathbf{u})$ for some $R \in \{R^C, R^V\}$ and some utility profile $\mathbf{u}$, where $b_r \neq a_r$ for some $r \in N$. It follows from Proposition 2 that $\mathbf{b}$ satisfies condition (1). If condition (2) is violated, voter r can increase her utility by ε, by changing her vote to a_r, as c_j would remain the unique election winner in this case. If condition (3) is violated for some $i \in N$ and some $c_k \in H(\mathbf{b})$, voter i can profitably deviate by changing her vote to c_k; if $b_i = c_j$, c_k would then become the unique election winner, and if $b_i \neq c_j$, the outcome will be a tie between c_j and c_k, so under R each of them will win with positive probability. Similarly, if condition (4) is violated for some $i \in N$ and some $c_\ell \in H'(\mathbf{b})$, voter i can profitably deviate by changing her vote to c_ℓ, so that the outcome becomes a tie among $H(\mathbf{b}) \cup \{c_j, c_\ell\}$. This concludes the proof. □

Corollary 3. *For $R \in \{R^C, R^V\}$, $(\mathcal{T}, R)$-*SINGLENE*, $(\mathcal{T}, R)$-*TIENE*, and $(\mathcal{T}, R)$-*EXISTNE *are* NP*-complete.*

Proof. Let $R \in \{R^C, R^V\}$. For $(\mathcal{T}, R)$-TIENE, as mentioned above, our claim follows from Theorem 2 and its implications, as discussed in the section on lazy voters.

For $(\mathcal{T}, R)$-SINGLENE with $R \in \{R^C, R^V\}$, we can use the same reduction from MSI as in Theorem 3. The only change is the analysis in the last part of the proof of Theorem 3, due to the different tie-breaking rule. In particular, suppose again that a PNE $\mathbf{b}$ exists, where w_2 is the winner. By the analysis of Theorem 3, the set of candidates $\mathcal{E}' = \{e \in \mathcal{E} \mid \mathrm{sc}(e, \mathbf{b}) = s - 2\}$ contains at least q elements. Consider a candidate $e \in \mathcal{E}'$. Suppose that among the voters from Block 1 who deviated to w_2, there exists a voter v_{i_j} who prefers e to w_2. Her utility in $\mathbf{b}$ is $1/2$. Suppose that she deviates to e instead. In this case the score of w_1, w_2, and e becomes $s - 1$. Therefore, the new winning set is $\{w_1, w_2, e\}$. Given that $u_{i_j}(w_1) = 1/4$ and $u_{i_j}(e) > 3/4$, the utility of voter v_{i_j} becomes more than $1/2$, contradicting the fact that $\mathbf{b}$ is a PNE. Thus, for any voter v_{i_j} in Block 1 who deviated to w_2, and for any candidate $e \in \mathcal{E}'$ it holds that v_{i_j} prefers w_2 to e, i.e., $e \in A_{i_j}$. Thus, $e \in A_{i_j}$ for each $j \in [k]$, where the sets $A_{i_1}, \ldots, A_{i_k}$ are defined in the proof of Theorem 3. As this holds for every $e \in \mathcal{E}'$ and $|\mathcal{E}'| \geq q$, this means that we have a "yes"-instance of MSI. For the reverse direction, the arguments are very similar to the reverse direction in the proof of Theorem 3.

Finally, regarding $(\mathcal{T}, R^L)$-EXISTNE, a simple modification in the reduction of Theorem 3 can yield the desired result; we omit the details from this version. □

C Price of Anarchy under Lexicographic Tie-breaking

We show that $\mathrm{PoA} = \Omega(n)$ both for lazy and for truth-biased voters under lexicographic tie-breaking. In particular, we first establish that $\mathrm{PoA} = n - 2$ for lazy voters. Then we show that $\mathrm{PoA} = 2n/3$ in the truth-biased model. Similar results can be proved for randomized tie-breaking, and we omit them from this version of the paper.

Proposition 4. *For lexicographic tie-breaking and lazy voters,* $\mathrm{PoA} = n - 2$.

Proof. We prove first that $\mathrm{PoA} \leq n-2$. To see this, note that by Theorem 1, the winner in any PNE must have a positive score in the truthful profile. Thus, in the worst-case

scenario for the Price of Anarchy, the truthful winner of $\mathbf{a}$ has score $n-1$, and there is a PNE where the winner is the candidate supported by the remaining voter. Thus $\mathrm{PoA} \leq n-2$.

To show the lower bound it suffices to exhibit an example. This is done in Example 2 below. □

Example 2. Consider the lazy voters model and the profile of Figure 1, with n voters and n candidates. It does not matter how we fill in the missing rankings in the figure. The truthful winner is c_3 with a score of $n-1$. However, consider the profile $\mathbf{b} = (c_2, \bot, \bot, \ldots, \bot)$. The winner in $\mathbf{b}$ is c_2, and no voter can unilaterally change the outcome in her favor. Indeed, if anyone votes for c_1, then c_1 is the new winner, but all voters prefer c_2 to c_1. On the other hand, voting for any other candidate cannot change the outcome due to tie-breaking. Since the score of c_2 in $\mathbf{a}$ is 1, we have $\mathrm{PoA} \geq n-2$.

1	2	3	...	n
c_2	c_3	c_3	...	c_3
⋮	c_2	c_2	...	c_2
⋮	c_1	c_1	...	c_1
⋮	⋮	⋮	...	⋮

Figure 1: PoA example for lazy voters.

Proposition 5. *For lexicographic tie-breaking and truth-biased voters,* $\mathrm{PoA} = 2n/3$. *This holds even for single-peaked or single-crossing preference profiles.*

Proof. As in Proposition 4, we first prove the upper bound. Let c_i be the winner in the truthful profile with a score of s^*. Let $\mathbf{b} \neq \mathbf{a}$ be a non-truthful PNE and let c_j be the winner in $\mathbf{b}$. Clearly, we have $\mathrm{PoA} \leq s^*$, since in the worst case c_j has no supporters in $\mathbf{a}$. Hence, it is enough to bound s^*.

By Lemma 1, we know that there exists at least one threshold candidate with respect to $\mathbf{b}$. We consider two cases:

Case 1: $c_i \notin T(\mathbf{b})$. Then there is some $c_k \neq c_i$ such that $c_k \in T(\mathbf{b})$ Let $s = \mathrm{sc}(c_k, \mathbf{a})$. By Lemma 1 we know that c_k receives s points in $\mathbf{b}$ as well. Hence c_j has a score of at most $s+1$ in $\mathbf{b}$. By Proposition 2 this means that there are at most $s+1$ non-truthful votes in $\mathbf{b}$. Hence the score of c_i in $\mathbf{b}$ has to be at least $s^* - (s+1)$. Since c_i is not a winner in $\mathbf{b}$, we have $s^* - (s+1) \leq \mathrm{sc}(c_i, \mathbf{b}) \leq s+1$, and hence $s^* \leq 2s+2$. Since the total score of c_i and c_k in $\mathbf{a}$ does not exceed n, we have $s + s^* \leq n$. But then, if $s^* > 2n/3$, this would imply that $s > n/3 - 1$, i.e., $s \geq n/3$, and hence $s + s^* > n$, a contradiction. Thus we have $\mathrm{PoA} \leq s^* \leq 2n/3$.

Case 2: $c_i \in T(\mathbf{b})$. In this case the Price of Anarchy is somewhat better. Let $s = \mathrm{sc}(c_j, \mathbf{b})$. Candidate c_i must have the same set of votes in $\mathbf{b}$ as in $\mathbf{a}$ by Lemma 1. Hence we have $s + s^* \leq n$. But we must also have $s^* \leq s$, otherwise c_j is not the winner. But then if $s^* > n/2$, we would also have $s > n/2$, a contradiction. Thus, in this case we have $\mathrm{PoA} \leq s^* \leq n/2$.

Hence in worst case, $\mathrm{PoA} \leq 2n/3$. Finally, to show that the worst case is attained, we exhibit a construction in Example 3. □

Example 3. In Figure 2, we show a preference profile for n voters, where n is divisible by 3. Block 1 consists of $n/3$ voters, Block 2 consists of $n/3+1$ voters, and Block 3 has $n/3-1$ voters. In the figure, it does not matter how we fill in the missing rankings, but note that we can fill them in a way that makes the preference profile single-peaked or single-crossing.

Suppose the tie-breaking rule is $c_1 > c_2 > c_3$. Under truthful voting, c_3 is the winner with a score of $2n/3$. We claim now that the profile $\mathbf{b}$, in which all voters of Block 2 vote for c_2 is a PNE. To see this, note that c_2 is indeed the winner in $\mathbf{b}$ with a score of $n/3+1$. Candidate c_1 would only need one additional vote to become the winner, but there is no incentive for any voter from Block 2 or 3 to vote for c_1, since all of them prefer c_2 to c_1. Also, no voter from Block 2 can change the outcome in favor of c_3 by a unilateral deviation, due to the tie-breaking rule. If a voter from Block 2 switches to her truthful vote, then the new winner is c_1, since there is a tie with all candidates. Hence $\mathbf{b}$ is a PNE, and the score of c_2 in the truthful profile is 0. This means that in this example we have $\mathrm{PoA} \geq 2n/3$.

Block 1				Block 2				Block 3			
c_1	c_1	...	c_1	c_3	c_3	...	c_3	c_3	c_3	...	c_3
⋮	⋮	...	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮
arbitrary				arbitrary				arbitrary			
⋮	⋮	...	⋮	⋮	⋮	...	⋮	⋮	⋮	...	⋮
c_2	c_2	...	c_2	c_2	c_2	...	c_2	c_2	c_2	...	c_2
c_3	c_3	...	c_3	c_1	c_1	...	c_1	c_1	c_1	...	c_1

Figure 2: PoA example for truth-biased voters

D Principled Voters

We will now present our results for the setting with principled voters. We omit the proofs of all results in this section, as they follow directly from the analysis presented earlier in the paper.

In what follows, we consider elections with a set of strategic (i.e., lazy or truth-biased) voters $N = \{1, \ldots, n\}$ and a set of principled voters $P = \{n+1, \ldots, n+s\}$; we assume that either all voters in N are lazy or all of them are truth-biased. We extend our notation to such games as follows: given a setting $\mathcal{S} \in \{\mathcal{L}, \mathcal{T}\}$, a tie-breaking rule $R \in \{R^L, R^V, R^C\}$, n strategic voters with utilities $\mathbf{u} = (u_1, \ldots, u_n)$, and s principled voters, whose votes are given by the ballot vector $\mathbf{a}^P = (a_{n+1}, \ldots, a_{n+s})$, we denote the resulting game by $(\mathcal{S}, R, \mathbf{u}, \mathbf{a}^P)$. The principled voters are not considered to be among the players; thus, the set of players in the modified game is still N. As before, we use $\mathbf{b}$ to denote a ballot vector of the strategic voters; $\mathbf{b}+\mathbf{a}^P$ denotes a ballot vector

that combines the votes of the strategic and principled voters. The computational problems EXISTNE, TIENE, and SINGLENE extend naturally to this setting; we denote the respective variants of these problems by EXISTNE^P, TIENE^P, and SINGLENE^P, respectively.

Principled + Lazy Voters, Lexicographic Tie-breaking

We have argued that in elections where all voters are lazy and the tie-breaking rule is R^L, there is at most one voter who does not abstain and all PNE have the same winner. However, in the presence of principled voters, this is no longer true; indeed, there are elections where *every* candidate can win in a PNE.

Example 4. Consider an election over a candidate set $C = \{c_1, \ldots, c_m\}$, $m > 1$, where there are two principled voters who both vote for c_m, and two lazy voters who both rank c_m last. Then the ballot vector where both lazy voters abstain is a PNE (with winner c_m). Moreover, for every $j \in [m-1]$ the ballot vector where both lazy voters vote for c_j is a PNE as well (with winner c_j).

Nevertheless, given an election with principled and lazy voters, we can characterize the set of candidates who can win in a PNE of the respective game.

Proposition 6. *Let $\mathbf{u}$ be the lazy voters' utility profile over C and let $\mathbf{a}^P$ be the principled voters' ballot vector. Let $j = \min\{k \mid c_k \in W(\mathbf{a}^P)\}$, and let $H^+(\mathbf{a}^P) = \{c_k \in H(\mathbf{a}^P) \mid k < j\}$. Then the game $G = (\mathcal{L}, R^L, \mathbf{u}, \mathbf{a}^P)$ has the following properties.*

(1) If $\mathbf{b}$ is a PNE of G then there is at most one candidate $c \in C$ such that $b_i = c$ for some $i \in N$; further, if $b_i = c$ for some $c \in C$, $i \in N$, then c is the winner in $\mathbf{b} + \mathbf{a}^P$.

(2) G has a PNE where c_j wins if and only if $(\bot, \ldots, \bot)$ is a PNE of G.

(3) If $k > j$ then G has a PNE where c_k wins if and only if there are at least $M(\mathbf{a}^P) + 1 - \mathrm{sc}(c_k, \mathbf{a}^P)$ lazy voters who prefer c_k to all candidates in $(W(\mathbf{a}^P) \cup H^+(\mathbf{a}^P)) \setminus \{c_k\}$.

(4) If $k < j$ then G has a PNE where c_k wins if and only if there are at least $M(\mathbf{a}^P) - \mathrm{sc}(c_k, \mathbf{a}^P)$ lazy voters who prefer c_k to all candidates in $(W(\mathbf{a}^P) \cup H^+(\mathbf{a}^P)) \setminus \{c_k\}$.

Corollary 4. *The problems $(\mathcal{L}, R^L)$-EXISTNE^P, $(\mathcal{L}, R^L)$-TIENE^P, and $(\mathcal{L}, R^L)$-SINGLENE^P are in* P.

Principled + Lazy Voters, Randomized Tie-breaking

We will now consider the effect of the presence of principled voters on lazy voters under randomized tie-breaking. We show that single-winner PNE in this setting may have a more complicated structure than single-winner PNE in the absence of principled voters. On the other hand, PNE where several candidates are tied for winning are very similar to those that arise when no principled voters are present. We first consider the random candidate tie-breaking rule.

Proposition 7. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be the lazy voters' utility profile over C, $|C| = m$, and let $\mathbf{a}^P = (a_{n+1}, \ldots, a_{n+s})$ be the principled voters' ballot profile. The game $G = (\mathcal{L}, R^C, \mathbf{u}, \mathbf{a}^P)$ admits a PNE $\mathbf{b}$ with $W(\mathbf{b} + \mathbf{a}^P) = \{c_j\}$ for some $c_j \in C$ if and only if one of the following conditions holds:*

(1) $W(\mathbf{a}^P) = \{c_j\}$, $H(\mathbf{a}^P) = \emptyset$;

(2) $|V_j| \geq M(\mathbf{a}^P) + 1 - \mathrm{sc}(c_j, \mathbf{a}^P)$, where V_j is the set that consists of all voters $i \in N$ such that (a) $u_i(c_j) > u_i(c_k)$ for all $c_k \in W(\mathbf{a}^P)$ and (b) for each $c_\ell \in H(\mathbf{a}^P)$ it holds that

$$u_i(c_j) \geq \frac{1}{|W(\mathbf{a}^P) + 1|} \sum_{c \in W(\mathbf{a}^P) \cup \{c_\ell\}} u_i(c).$$

Moreover, if condition (1) holds then G has a PNE where all lazy voters abstain, and if condition (2) holds then G has a PNE where exactly $M(\mathbf{a}^P) + 1 - \mathrm{sc}(c_j, \mathbf{a}^P)$ lazy voters vote for c_j, while the remaining lazy voters abstain. The game G has no other PNE with winning set $\{c_j\}$.

Corollary 5. *The problem $(\mathcal{L}, R^C)$-SINGLENE^P is in* P.

Proposition 8. *Let $\mathbf{u} = (u_1, \ldots, u_n)$ be the lazy voters' utility profile over a candidate set C, $|C| = m$, and let $\mathbf{a}^P = (a_{n+1}, \ldots, a_{n+s})$ be the principled voters' ballot profile. Then the game $G = (\mathcal{L}, R^C, \mathbf{u}, \mathbf{a}^P)$ admits a PNE $\mathbf{b}$ with $|W(\mathbf{b} + \mathbf{a}^P)| > 1$ if and only if one of the following conditions holds:*

(1) each candidate is ranked first by at most one voter in $N \cup P$ and $\frac{1}{n+s} \sum_{i \in N \cup P} u_\ell(a_i) \geq \max_{i \in (N \cup P) \setminus \{\ell\}} u_\ell(a_i)$ for all $\ell \in N$.

(2) there exists a set of candidates $X = \{c_{\ell_1}, \ldots, c_{\ell_k}\}$ with $k \geq 2$, a positive integer $n' \leq n$ with $n'/k \geq 2$ such that for each $c \notin X$ we have $\mathrm{sc}(c, \mathbf{a}^P) < n'/k$, and a partition of the lazy voters into k groups $N_1, \ldots, N_k$ (some of which may be empty) such that

(a) for each $j \in [k]$ we have $|N_j| + \mathrm{sc}(c_{\ell_j}, \mathbf{a}^P) = n'/k$;

(b) for each $j \in [k]$ and each $i \in N_j$ we have $c_{\ell_j} \succ_i c$ for all $c \in X \setminus \{c_{\ell_j}\}$;

(c) for each $j \in [k]$ and each $i \in N_j$ we have $\frac{1}{k} \sum_{c \in X} u_i(c) \geq \max_{c \in X \setminus \{c_{\ell_j}\}} u_i(c)$;

(d) for each $j \in [k]$, each $i \in N_j$, and each $c' \in C \setminus X$ with $\mathrm{sc}(c', \mathbf{a}^P) = n'/k - 1$ we have $\frac{1}{k} \sum_{c \in X} u_i(c) \geq \frac{1}{k} \sum_{c \in (X \cup \{c'\}) \setminus \{c_{\ell_j}\}} u_i(c)$.

Moreover, if condition (1) holds then G has a PNE where each lazy voter votes for her top candidate, and if condition (2) holds, then G has a PNE where each lazy voter votes for her top candidate in X. The game G has no other PNE with two or more winners.

The following corollary is a direct consequence of Corollary 2 and the fact that the model with no principled voters is a special case of the model with principled voters.

Corollary 6. *The problems $(\mathcal{L}, R^C)$-TIENE^P and $(\mathcal{L}, R^C)$-EXISTNE^P are* NP*-complete.*

The reader may have noticed that Proposition 7, Corollary 5, and Proposition 8 are stated for R^C, but not for R^V. The reason for this is that in the presence of principled voters the tie-breaking rules R^C and R^V are no longer equivalent.

Example 5. Consider an election over the candidate set $C = \{c_1, c_2, c_3\}$, where there are two lazy voters whose utility function is given by $u(c_1) = 20$, $u(c_2) = 4$, $u(c_3) = 1$, two lazy voters whose utility function is given by $u'(c_1) = 20$, $u'(c_2) = 4$, $u'(c_3) = 1$, and one principled voter who ranks the candidates as $c_3 \succ c_1 \succ c_2$. It is easy to see that both for R^C and for R^V the resulting game has a PNE where two lazy voters vote for c_1, two lazy voters vote for c_2, and the principled voter votes for c_3. Under R^C candidates c_1 and c_2 are equally likely to win in this PNE. However, under R^V candidate c_1 wins with probability $3/5$ and candidate c_2 wins with probability $2/5$.

Nevertheless, all results in this section can be extended to random voter tie-breaking, by replacing the uniform lotteries over the winning sets in Propositions 7 and 8 by lotteries that correspond to choosing an element of the winning set according to the preferences of a random voter (who may be principled or lazy). While the inequalities that one needs to verify become more cumbersome, the complexity of the respective computational problems remains the same. In particular, we obtain the following corollary.

Corollary 7. *The problems* $(\mathcal{L}, R^V)$-TIENE^P *and* $(\mathcal{L}, R^V)$-EXISTNE^P *are* NP*-complete, whereas* $(\mathcal{L}, R^V)$-SINGLENE^P *is in* P.

Principled + Truth-biased Voters

Principled and truth-biased voters are quite similar in their behavior; therefore, adding principled voters to the setting of Section 4 results in fewer changes than adding them to the setting of Section 3.

To illustrate this point, we will now show how to extend Proposition 3 to settings where principled voters may be present.

Proposition 9. *Let* $\mathbf{u}$ *be the utility profile of truth-biased voters, let* $\mathbf{a}$ *be their truthful ballot vector, and let* $\mathbf{a}^P$ *be the ballot vector of principled voters. Let* $j = \min\{r \mid c_r \in W(\mathbf{a} + \mathbf{a}^P)\}$. *Then* $\mathbf{a}$ *is a PNE of* $(\mathcal{T}, R^L, \mathbf{u}, \mathbf{a}^P)$ *if and only if neither of the following conditions holds:*

(1) $|W(\mathbf{a} + \mathbf{a}^P)| > 1$, *and there exists a candidate* $c_k \in W(\mathbf{a} + \mathbf{a}^P)$ *and a voter* $i \in N$ *such that* $a_i \neq c_k$ *and* $c_k \succ_i c_j$.

(2) $H(\mathbf{a}+\mathbf{a}^P) \neq \emptyset$, *and there exists a candidate* $c_k \in H(\mathbf{a}+\mathbf{a}^P)$ *and a voter* $i \in N$ *such that* $a_i \neq c_k$, $c_k \succ_i c_j$, *and* $k < j$.

All other claims in Section 4 can be modified in a similar way: essentially, we replace $W(\mathbf{b})$, $H(\mathbf{b})$ and $H'(\mathbf{b})$ with $W(\mathbf{b} + \mathbf{a}^P)$, $H(\mathbf{b} + \mathbf{a}^P)$ and $H'(\mathbf{b} + \mathbf{a}^P)$, but when considering the voters' incentives to change their votes, we limit our attention to truth-biased voters. Of course, we have to take into account the number of votes cast by principled voters in favor of each candidate in the winning set, and distinguish between R^V and R^C (as we did above). Finally, it is immediate that all hardness results established in Section 4 remain true in the presence of principled voters.